



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 17/30	A1	(11) International Publication Number: WO 99/32999 (43) International Publication Date: 1 July 1999 (01.07.99)
--	-----------	--

(21) International Application Number: PCT/US98/27240

(22) International Filing Date: 21 December 1998 (21.12.98)

(30) Priority Data:

60/068,415	22 December 1997 (22.12.97)	US
09/106,212	29 June 1998 (29.06.98)	US

(71)(72) Applicant and Inventor: NG, Tony, Chun, Tung
[CN/US]; 3716 Harlequin Terrace, Fremont, CA 94555 (US).

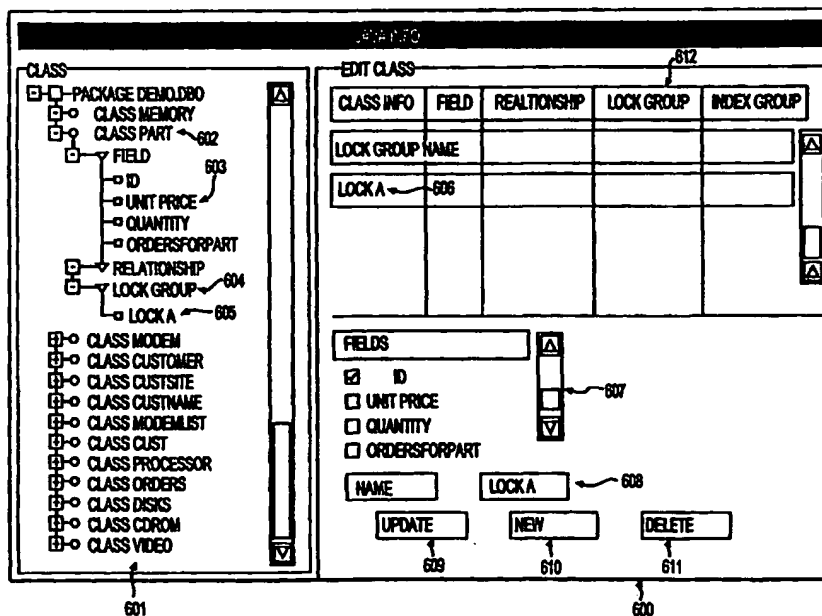
(74) Agents: GARRETT, Arthur, S. et al.; Finnegan, Henderson, Farabow, Garrett & Dunner, L.L.P., 1300 I Street, N.W., Washington, DC 20005-3315 (US).

(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published

*With international search report.**Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.*

(54) Title: USER INTERFACE FOR THE SPECIFICATION OF LOCK GROUPS



(57) Abstract

A user interface permits a programmer or other person to manage lock groups for classes. The programmer enters information through the user interface to define new lock groups, update defined lock groups, and delete lock groups. The programmer manages the lock groups in the classes, and an optional mapping tool maps the defined lock groups when converting data between an object model and a relational model.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

USER INTERFACE FOR THE SPECIFICATION OF LOCK GROUPS

REFERENCE TO RELATED APPLICATIONS

The following identified U.S. patent applications are relied upon and are incorporated in their entirety by reference in this application as if fully set forth.

Provisional U.S. Patent Application No. 60/068,415, entitled "System and Method for Mapping Between Objects and Databases," filed on December 22, 1997.

U.S. Patent Application No. _____, entitled "Object Relational Mapping Tool That Processes Views," bearing attorney docket no. 06502.0136-00000, and filed on the same date herewith.

U.S. Patent Application No. _____, entitled "Evolution Of Object-Relational Mapping Through Source Code Merging," bearing attorney docket no. 06502.0137-00000, and filed on the same date herewith.

U.S. Patent Application No. _____, entitled "Integrating Both Modifications to an Object Model and Modifications to a Database into Source Code by an Object-Relational Mapping Tool," bearing attorney docket no. 06502.0138-00000, and filed on the same date herewith.

U.S. Patent Application No. _____, entitled "User Interface for the Specification of Lock Groups," bearing attorney docket no. 06502.0142-00000, and filed on the same date herewith.

U.S. Patent Application No. _____, entitled "A Fine-Grained Consistency Mechanism for Optimistic Concurrency Control Using Lock Groups," bearing attorney docket no. 06502.0143-00000, and filed on the same date herewith.

U.S. Patent Application No. _____, entitled "User Interface for the Specification of Index Groups Over Classes," bearing attorney docket no. 06502.0144-00000, and filed on the same date herewith.

U.S. Patent Application No. _____, entitled "Method and Apparatus for Creating Indexes in a Relational Database Corresponding to Classes in an Object-Oriented Application," bearing attorney docket no. 06502.0145-00000, and filed on the same date herewith.

U.S. Patent Application No. _____, entitled "Method and Apparatus for Loading Stored Procedures in a Database Corresponding to Object-Oriented Data

Dependencies," bearing attorney docket no. 06502.0146-00000, and filed on the same date herewith.

U.S. Patent Application No. _____, entitled "An Integrated Graphical User Interface Method and Apparatus for Mapping Between Objects and Databases," bearing attorney docket no. 06502.0147-00000, and filed on the same date herewith.

U.S. Patent Application No. _____, entitled "Methods and Apparatus for Efficiently Splitting Query Execution Across Client and Server in an Object-Relational Mapping," bearing attorney docket no. 06502.0148-00000, and filed on the same date herewith.

FIELD OF THE INVENTION

The present invention relates to a user interface for specifying and managing lock groups for data stored in a database.

BACKGROUND OF THE INVENTION

Object-relational mapping tools facilitate development of application programs that utilize a relational database. A relational database stores data in tables having rows (records) and columns (fields). The tables are usually interrelated, and thus, there is a logical structure imposed on the database. This logical structure is known as a schema. Each table has a primary key, comprising one or more columns that uniquely identify a row. For example, in a table with rows of customers, a column storing each customer's social security number may be used as the primary key because it uniquely identifies each customer in the table. A table may also have another key, known as a foreign key, associating a row in one table to one or more rows in another table. For example, where one table contains customer information and another table contains order information for the customers, a foreign key may exist to relate one customer (or row) in the customer table with one or more orders (or rows) in the order table.

Object-relational mapping tools read a database and automatically generate source code from the database. This source code contains a number of classes whose interrelationships reflect the logical structure, or schema, of the database. A class, such as a Java™ class, is a data structure containing both data members that store data and function members (or methods) that act upon the data. The source code contains one class for each table in the database, and each class has a data member for each column in the corresponding

table. Additionally, the classes contain function members that are used to both get and set the values for the data members and, eventually, update the database.

By using an object-relational mapping tool, a programmer can automatically generate source code to facilitate their database application development. After the source code is generated, the programmer writes code to interact with the classes in the source code and not the database, thus hiding the complexities of interacting with the database from the programmer. This allows a programmer who is familiar with object-oriented programming to code against familiar classes and not unfamiliar, sometimes cumbersome to use, database query languages.

Conventional object-relational mapping tools, however, do not enable users to specify lock groups based on classes mapped from tables in a relational database. A need therefore exists for a mapping tool with this capability.

BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings, which are incorporated in and constitute a part of this specification, illustrate an implementation of the invention and, together with the description, serve to explain the advantages and principles of the invention. In the drawings,

Figure 1 depicts a data processing system suitable for practicing methods and systems consistent with the present invention;

Figure 2 depicts a more detailed diagram of the database depicted in Figure 1;

Figure 3 depicts a database data structure reflecting the schema of the database depicted in Figure 1;

Figure 4A depicts an object model containing information derived from the database data structure depicted in Figure 3;

Figure 4B depicts source code generated from the object model depicted in Figure 4A;

Figure 5 is a diagram of an exemplary lock group for an object in a class;

Figure 6 is an example of an user interface for defining lock groups for classes mapped from tables;

Figure 7 is a block diagram of software modules for accessing and operating on data using lock groups;

Figure 8 is a diagram of software modules for processing transactions on data using associated lock groups;

Figure 9 is a flow chart illustrating exemplary processing associated with the user interface shown in Figure 6; and

Figure 10 is a flow chart illustrating exemplary processing associated with applications performing transactions on data in a database.

SUMMARY OF THE INVENTION

Apparatus and methods consistent with the present invention provide a user interface to allow a programmer or other person to view and enter information relating to lock groups, potentially for use with a mapping tool that maps corresponding data between a relational model and an object model.

A method consistent with the present invention provides a user interface having information representing lock groups for a class structure. Information representing one or more of the lock groups is selectively displayed through the user interface.

Another method consistent with the present invention provides a user interface having information representing data in a database and receives through the user interface information for defining a lock group associated with the data. The defined lock group is stored for use in processing associated with the data.

An apparatus consistent with the present invention provides a user interface having information representing lock groups for a class structure. The apparatus selectively displays through the user interface information representing one or more of the lock groups.

Another apparatus consistent with the present invention provides a user interface having information representing data in a database and receives through the user interface information for defining a lock group associated with the data. The apparatus stores the defined lock group for use in processing associated with the data.

DETAILED DESCRIPTION

Methods and systems consistent with the present invention provide a user interface for a programmer or other person to specify lock groups among classes of objects potentially mapped from a relational database. The interface may also be used to modify or delete existing lock groups. A mapping tool can map the lock groups during mapping of data between objects in an object-oriented model and tables in a relational model and, therefore, a user's specified lock groups are saved and need not be repeatedly recreated.

Overview of Mapping Tool

In accordance with methods and systems consistent with the present invention, the improved object-relational mapping tool maps a database by first querying the database to determine its schema and then by creating an internal data structure (known as the "database data structure") representing that schema. From this data structure, the object-relational mapping tool creates an object model containing all of the information necessary to generate classes and then creates source code containing a number of Java classes that may be used by a programmer to interface with the database.

Implementation Details

Figure 1 depicts a data processing system 100 suitable for practicing methods and systems consistent with the present invention. Data processing system 100 includes computer 101 connected to a network 102 such as the Internet. Computer 101 includes memory 104, secondary storage device 106, central processing unit (CPU) 108, input device 110, and video display 112. Memory 104 includes an object-relational mapping tool 114 (ORMT) in accordance with methods and systems consistent with the present invention. In turn, the object-relational mapping tool 114 includes object model 116 and database data structure 115, reflecting the schema of database 118, stored on secondary storage device 106. Also stored on secondary storage device 106 is source code 120, containing classes reflecting the schema of database 118 and containing any customizations of the programmer.

Memory 104 also includes a runtime system 123, which includes a virtual machine (VM) 124. Secondary storage device 106 further contains a program 125 with source code and various class files 126. An exemplary runtime system for purposes of implementing methods and systems consistent with the principles of the present invention includes the Java™ runtime system included in the Java™ Development Kit from Sun Microsystems, Inc. The Java runtime system includes a Java VM. The Java VM is described in Lindholm and Yellin, The Java™ Virtual Machine Specification, Addison-Wesley (1997), which is incorporated herein by reference.

Although computer 101 is depicted with various components, one skilled in the art will appreciate that this computer can contain additional or different components. Additionally, although computer 101 is shown connected to network 102, computer 101 may be connected to other networks, including other wide area networks or local area networks.

Furthermore, although aspects of the present invention are described as being stored in memory, one skilled in the art will appreciate that these aspects can also be stored on or read from other types of computer program products or computer-readable media, such as secondary storage devices, including hard disks, floppy disks, or CD-ROM; a carrier wave from the Internet; or other forms of RAM or ROM. Still further, one skilled in the art will appreciate that database 118 and source code 120 may be stored on or distributed across other devices on network 102. In addition, the computer-readable media may include instructions for controlling a computer systems, such as computer 101, to perform a particular method.

Sun, Sun Microsystems, the Sun logo, Java™, and Java™-based trademarks are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

Object-relational mapping tool 114 reads database 118 to examine its schema, constructs database data structure 115 to reflect this schema, generates an object model 116 based on database data structure 115, and then creates source code 120 based on object model 116. It should be noted that, at the time object model 116 is generated, the object-relational mapping tool allows the programmer to add customizations, and these customizations will be reflected in the source code 120. For example, the programmer may add a new method, rename a field (and use it for a different purpose), change the attributes of a field (e.g., the type or whether it can accept a null value), or override the mapping of a field. When a field mapping is overridden, that field will not appear in the source code.

Figure 2 depicts a more detailed diagram of an example of database 118, containing a customer table 202 and an order table 204. The customer table 202 includes a customer ID column 206, and a name column 208. The customer ID column 206 serves as the primary key for the customer table 202. The order table 204 includes order ID column 212, date column 214, and customer ID column 216. The order ID column 212 serves as the primary key for the order table 204. Customer ID column 216 is the foreign key to customer ID column 206, meaning customer ID column 216 refers to the customer ID column 206 in one or more rows. As previously stated, database data structure 115 represents the schema of database 118. Object-relational mapping tool 114 creates database data structure 115 by querying database 118 to identify its schema and by creating the data structure to reflect the schema. This process is known as "importing" the database schema and is described in

further detail below. Once created, database data structure 115 appears as shown in Figure 3 and includes an object 302, reflecting the customer table 202, and an object 304, reflecting the order table 204. Object 302 contains a list 306 of foreign key objects, if any, each containing the name of the foreign key as well as an indication of the columns that comprise the foreign key. Additionally, object 302 contains a list 308 of the indexes in the customer table 202, where each element of the list is an index object containing an indication of the type of index (e.g., primary, non-unique, and unique) and a list of columns that comprise the index. Object 302 also contains a hash table 310, where each entry in the hash table is a column object 312, 314 containing data for a particular field, including its name, type, and length. Object 304 contains similar information, including a list of foreign keys 316, a list of indexes 318, and a hash table 320 with column objects 322-326 for each field or column.

Using the object-relational mapping tool, the programmer may customize the object model. For example, the programmer may rename the name field to SSN and may subsequently use this field to store the customer's social security number, in which case the customer's social security number will be stored in the name field 208 of the database 118. By making such a customization, it is reflected in the object model 116 shown in Figure 4A. Object model 116, generated by the object-relational mapping tool, contains the programmer's customization (e.g., the name field has been renamed to SSN). Object model 116 contains objects 400 and 401, representing an intermediate form of the information for a class before it is written as source code. Object 400 contains information for the customer table 202, including a list 402 of relationship objects 403, each containing information indicating a relationship (i.e., a foreign key). For example, relationship object 403 links a field in object 400 with a field in object 401. Additionally, object 400 contains a hash table 404 with an entry 406, 408 for each field in customer table 202, each entry containing the name and type of the field. Similarly, object 401 contains information for order table 204, including a list 410 of relationship objects 411 and a hash table 412 containing entries 414-418 for each field in order table 204.

As can be appreciated from this description of object model 116, it contains all of the information necessary to create the classes in the source code, an example of which is depicted in Figure 4B. Figure 4B depicts source code file 116 with the Java™ programming language representation of objects 400 and 401. Class 420 reflects customer table 202 and

class 424 reflects order table 204. As such, class 420 contains a data member for customer ID, social security number, and a collection of objects representing the orders associated with that particular customer, thus implementing the foreign key. Class 420 also contains a number of methods to both get and set the value of the data members, including an iterator method to iterate through the order for this particular customer. Class 424 includes data members order ID and date and also includes various methods to both set and get the values for these data members. Additionally, class 424 contains a field, `Customer_for_Order`, implementing the foreign key with a reference to the particular customer object that placed that order.

When a foreign key is contained in the object model, the object-relational mapping tool typically creates a relationship in the source code between two classes to implement the foreign key. As stated above, with a foreign key, one or more records in one table (the referring table) refers to one record in another table (the referred table). This relationship is a one-to-many relationship, although it may be a one-to-one relationship. Additionally, instead of being bidirectional, the relationship may be unidirectional. To define this relationship in the Java™ programming language, the class representing the referring table is defined to have a member that is a collection of the class representing the referred table. A "collection" refers to a type indicating a grouping of instances of other classes. Then, in the class reflecting the referred table, a member is added providing a reference to the class that refers to it. For most cases, this is how a foreign key is implemented. However, when the foreign key for two tables overlaps with the primary key for those tables, it is more efficient to simply subclass the class reflecting the referred class from the class reflecting the referring class.

Lock Groups

A lock group specifies one or more fields of a class for an object model of data stored in a database for use by a system in avoiding conflicts that may arise when updating the database. The object model may be mapped from a relational model of the data. Figure 5 shows an example of a customer object 500 in an object model having a name field 501 and an address field 502. An exemplary lock group 503 is shown as including name field 501 and address field 502. With this lock group defined, only one user may modify data in the name and address fields during a transaction.

A system managing a database for which object 500 is defined uses exemplary lock group 503 to determine conflicts among applications processing data in the database. For example, when an application performs a transaction modifying data in object 500, including data in fields 501 or 502, and attempts to commit the transaction to save the modifications, a system determines if any other application modified data in fields 501 and 502. The system does not commit the transaction and modify the database if another application made modifications.

Lock groups are applicable, for example, to optimistic concurrency control with respect to applications updating a database. Under optimistic concurrency control, a system does not prevent access to fields in object; in other words, it does not lock out an application from accessing fields while another application may access those same fields. Rather, it typically permits applications to access any field in objects and uses lock groups to determine conflicts when an application commits data changes. Conflicts involve two or more applications attempting to modify and commit changes to the same field or set of fields. In comparison, under pessimistic concurrency control, a system avoids conflicts by locking out particular fields or a set of fields when an application accesses them, prohibiting other applications from updating the locked fields.

Lock group 503 is only one example of a lock group for fields in a class corresponding to an object model. Lock groups may typically include any number of fields in an object or fields among multiple objects. A default lock group will typically include all fields within an object. Therefore, allowing a user to define lock groups as fewer than all fields permits more narrowly defining what constitutes a conflict, permitting simultaneous access to data in an object and hence increasing data processing capability.

User Interface for Managing Lock Groups

An implementation consistent with the present invention relates to a user interface for specifying and managing lock groups within a database. A "user interface" refers to a mechanism to view a representation of information in a computer system, such as computer 101, and to enter information into the computer system. The input information may include commands instructing the computer to perform a particular function.

Figure 6 illustrates an example of such a user interface 600. Video display 112 may present user interface 600 in order to permit a user to view information relating to lock

groups, and a user may use input device 110, such as a cursor-control device, to enter or modify information through user interface 600. Lock groups defined through user interface 600 are typically stored in computer 101 such as in memory 104 or secondary storage 106. The appearance and structure of user interface 600 are only one such example of a user interface for accomplishing functions relating to lock groups, and implementations consistent with the present invention include any mechanism to display and enter information relating to lock groups.

When a user selects lock group button 612, the user may enter and define lock groups for particular classes, shown in a section 601. A "section" refers to an area of a user interface in which a system presents information or receives information from a user. The act of selecting may include using a cursor-control device to "click on" or "double click on" a particular item in the user interface. User interface 600 includes section 601 in order to illustrate fields of a particular class representing data in a database. A user may select the boxes next to the classes in order to expand the listing and view information related to each class, including fields of the class and lock groups defined for the class. For example, a class Part 602 is shown as having a number of fields 603 and as associated with one or more lock groups 604. In this example, one lock group 605 is defined for class Part, lock A 605. When a user selects lock A 605, a section 606 presents the corresponding lock group and a section 607 presents the fields of class Part 602. A user may select other classes in section 601 in a similar manner, and fields for those classes are also presented in section 607 when corresponding classes are selected. Sections 606 and 607 include scroll bars so that a user may scroll through lock groups or fields if there is insufficient room to illustrate all the lock groups or fields within the section.

A user may perform functions on lock groups by selecting buttons 609-611. In particular, by selecting update button 609, a user may update a defined lock group, and by selecting new button 610, a user may define and enter a new lock group. In order to update a lock group or define a new lock group, a user selects the boxes next to the particular fields in section 607 in order to select or deselect those fields to include within the lock group shown in section 606. When a user selects or "clicks on" a field or the box next to it, a check mark appears in the box next to the field in section 607 to indicate its selection. When a user "clicks on" the field again, or the box next to it, to deselect the field, the check mark

disappears, indicating deselection of the corresponding field. Other types of indicators are possible for showing selection of a field, such as highlighting the field, displaying the field within a box or in a different color, or displaying some other type of icon next to the field.

In section 608, a user may enter a name of a new lock group or modify the name of a defined lock group, and the entered name is shown in section 606. When a user selects update button 609, the system saves the new lock group or updates a defined lock group according to the user's selections, and the new or updated lock group is saved in the system and presented within section 601 for the corresponding class.

By selecting delete button 611, a user may delete a defined lock group. In particular, a user first selects a lock group in section 601 and then selects delete button 611, after which the system deletes that defined lock group and removes it from section 601.

Lock Group Modules

Figure 7 is a block diagram of software modules for implementing the exemplary user interface shown in Figure 6. These modules may operate within the hardware elements shown in Figure 1. In system 700, user interface 600 is typically part of a mapping tool 701 such as ORMT 114. Mapping tool 701 includes a relational model 702, an object model 704, and a mapping engine 703 for converting data between the relational model and the object model. Relational model 702 stores a representation of data from database 707 in tables, and object model 704 stores a representation of the data in objects. Relational model 702 may correspond to database data structure 115, and object model 704 may correspond to object model 116. Mapping tool 701 interacts with VM 705, which may correspond to VM 124. System 700 is connected to a database 707 for storing data in relational form, and database 707 may correspond to database 118.

A user operating with user interface 600 interacts with mapping tool 701 in order to enter and define the lock groups. Mapping engine 703 converts data between relational model 702 and object model 704, and may perform the mapping using the exemplary mapping tool described above. In performing the conversions, mapping engine 703 uses lock groups defined through interface 600 in order to map the lock groups between the relational and object models. The user-entered definitions of lock groups may be saved, such as in memory 104 or secondary storage 106, and the system may associate the saved lock groups with the corresponding classes according to known methods. Furthermore, user interface 600

need not necessarily be used with a mapping tool and may, alternatively, be used in general to specify lock groups on a database.

Figure 8 is a block diagram of software modules including a runtime illustrating the operation of applications on data within a database. These modules may operate within the hardware elements shown in Figure 1. System 800 includes one or more applications 801 performing transactions on data within database 707. Application 801 interacts with a Java runtime 802, which may correspond to runtime 123 and includes a runtime cache 803 storing one or more objects 804 and 805 and a transaction manager 806. Runtime 802 also interacts with VM 705. System 800 may also access data in database 707.

Objects are generated in runtime cache 803 as required by application 801, which accesses and operates on data within the objects. When application 801 completes its transaction, it commits the transaction to transaction manager 806, which determines any conflicts. Transaction manager 806 tracks values of data elements in the database over time in order to determine the conflicts.

For example, Table 1 illustrates how transaction manager 806 may track values of data elements to determine conflicts. In this example, application 1 attempts to modify the value of a data element x. At time t1, the value of data element x is 5 and application 1 begins a transaction. At time t2, the value of data element x remains the same, but at time t3 application 2 begins a transaction and changes the value of data element x to 6. Application 2 commits the data at time t4, at which time the value of data element x is changed to 6 in the database, assuming no conflicts exist. At time t5, application 1 changes the value of data element x to 7 and attempts to commit the transaction.

Table 1		
<u>time</u>	<u>application 1</u>	<u>application 2</u>
t1	x=5	
t2		x=5
t3		x→6
t4		commit data
t5	x→7 (conflict)	

Transaction manager 806 determines if a conflict exists by comparing before and after values of data element x. Transaction manager 806 recorded the value of data element x (=5)

when application 1 began its transaction, and the system knows that application 1 intends to commit a value of 7 for data element x. Therefore, for no conflicts to exist, the before and after values of data element x must be 5 and 7. However, data element x now has a value of 6 in the database, and the before and after values are 6 and 7. Transaction manager 806 thus determines that the before and after values are incorrect when application 1 commits the transaction and it does not make the requested change to the value of data element x. The system thus records and tracks values of data elements when transactions begin in order to determine conflicts when application commit transactions, and other methods of determining conflicts are possible.

Lock Group Processing

Figure 9 is a flow chart illustrating exemplary processing associated with user interface 600 shown in Figure 6. This processing may be implemented by the modules shown in Figure 7 operating within the hardware elements shown in Figure 1. In process 900, a user first selects a lock group button 612 (state 901). The user then selects a class and optionally selects a lock group in section 601 (state 902). The user may select the class and lock group as shown in section 601 by selecting the appropriate box next to the class names. The user selects one of the buttons 609-611 in order to select a particular function (state 903).

If the user selected new button 610, the user enters the lock group name in section 608 and selects fields for the lock group in section 607 by selecting the boxes next to the field names (state 904). The user selects update button 609 in order to enter the defined new lock group (state 905), after which the system determines if at least one field is selected for the lock group (state 906). If no fields are selected, the system provides an error message (state

907), which may include presenting a section with a message indicating that at least one field must be selected. Otherwise, the system updates the lock group definition with the new information, stores the defined lock group, and includes a listing of it in section 601 (state 908).

To update a lock group, a user selects fields for the lock group using section 607 (state 909). When the user selects the update button, the lock group is updated in accordance with the field selection, if at least one field was selected (states 905-908). If the user selects delete button 611, the system deletes the lock group selected in state 902 (state 910) and shown in section 606.

Figure 10 is a flow chart of processing associated with the modules shown in Figure 8 and involving processing by transaction manager 806. This processing may be implemented by the software modules shown in Figure 8 operating within the hardware elements shown in Figure 1. In process 1000, application 801 initiates a transaction (state 1001). Application 801 reads objects from runtime cache 803 and modifies a set of fields for the transaction in one or more of the objects (state 1002). Application 801 commits the transaction by interacting with transaction manager 806 (state 1003). As identified above, committing a transaction involves the application requesting that the system change data in a database according to modifications made by the application.

For each object application 801 modified, transaction manager 806 determines the fields modified and the relevant lock groups involved (state 1004). Transaction manager 806 determines if a conflict exists by determining if any other application modified a field within the relevant lock groups (state 1005), which it may accomplish by tracking data values as described above. If a conflict exists, transaction manager 806 aborts the transaction and

typically provides an indication that the transaction was aborted, for example, by throwing an exception (state 1006), which may include informing the application that a conflict occurred and the requested changes were not entered. Otherwise, transaction manager 806 commits the transaction (state 1007) by updating the values of data elements in the database according to the changes made by application 801.

While the present invention has been described in connection with a preferred embodiment, many modifications will be readily apparent to those skilled in the art, and this application is intended to cover any adaptations or variations thereof. For example, other types of user interfaces and hardware for presenting the user interface, and other types of programming languages for implementing an embodiment consistent with the present invention, may be used without departing from the scope of the invention. This invention should be limited only by the claims and equivalents thereof.

WHAT IS CLAIMED IS:

1. A method of permitting a user to view lock groups for a class structure corresponding to data stored in a database, comprising:
 - providing a user interface having information representing the lock groups; and
 - selectively displaying information representing one or more of the lock groups.
2. A method of defining a lock group for a class structure corresponding to data stored in a database, comprising:
 - providing a user interface having information representing the data; and
 - receiving information for defining a lock group associated with the data.
3. The method of claim 2 wherein receiving information for defining a lock group includes
 - receiving a request to update the defined lock group.
4. The method of claim 2 wherein receiving information for defining a lock group includes
 - receiving a request to delete the defined lock group.
5. The method of claim 2 wherein receiving information for defining a lock group includes
 - receiving an identification of particular fields of the data to include within the defined lock group.

6. An apparatus, comprising:

a memory having program instructions for permitting a user to view lock groups for a class structure corresponding to data stored in a database; and

a processor, responsive to the program instructions, to

provide a user interface having information representing the lock groups, and

selectively display information representing one or more of the lock groups.

7. An apparatus, comprising:

a memory having program instructions for defining a lock group for a class structure corresponding to data stored in a database; and

a processor, responsive to the program instructions, to

provide a user interface having information representing the data in the

database, and

receive information for defining a lock group associated with the data.

8. The apparatus of claim 7 wherein receiving information for defining a lock group includes

receiving a request to update the defined lock group.

9. The apparatus of claim 7 wherein receive information for defining a lock group includes

receiving a request to delete the defined lock group.

10. The apparatus of claim 7 wherein receive information for defining a lock group includes

receiving an identification of particular fields of the data to include within the defined lock group.

11. A computer-readable medium containing instructions for controlling a computer system to perform a method for permitting a user to view lock groups for a class structure corresponding to data stored in a database, the method comprising:

providing a user interface having information representing the lock groups; and

selectively displaying through the user interface information representing one or more of the lock groups.

12. A computer-readable medium containing instructions for controlling a computer system to perform a method for use in defining a lock group for a class structure corresponding to data stored in a database, the method comprising:

providing a user interface having information representing the data in the database;

and

receiving information for defining a lock group associated with the data.

13. The computer-readable medium of claim 12 wherein receiving information for defining a lock group includes

receiving a request to update the defined lock group.

14. The computer-readable medium of claim 12 wherein receiving information for defining a lock group includes

receiving a request to delete the defined lock group.

15. The computer-readable medium of claim 12 wherein receiving information for defining a lock group includes

receiving an identification of particular fields of the data to include within the defined lock group.

16. An apparatus for permitting a user to browse lock groups for a class structure corresponding to data stored in a database, comprising:

means for providing a user interface having information representing the lock groups;

and

means for selectively displaying through the user interface information representing one or more of the lock groups.




FIG. 1

2/9

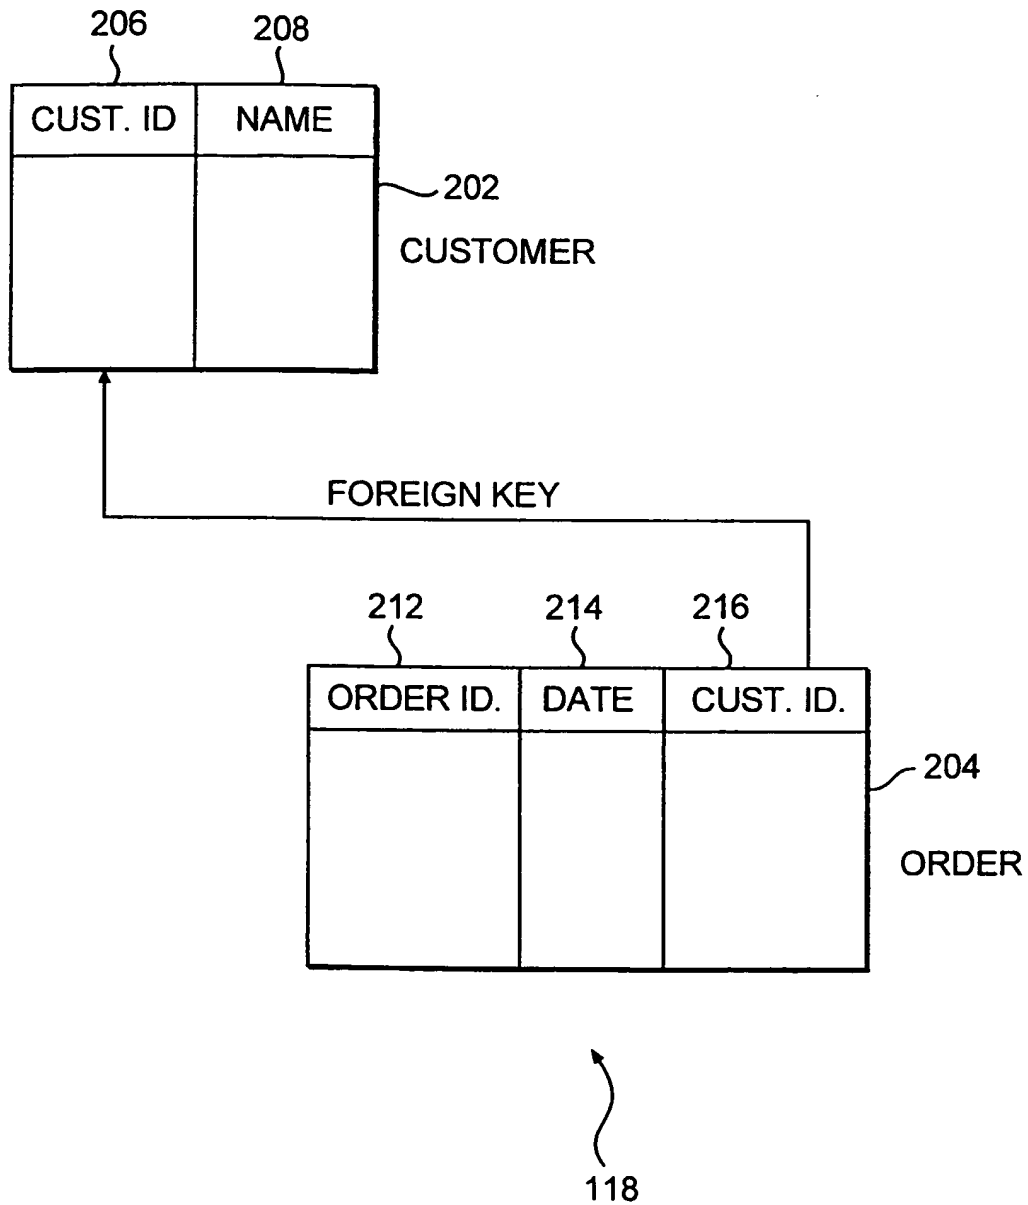


FIG. 2

3/9




FIG. 3




FIG. 4A

420

```

CLASS CUSTOMER {
    INT CUST_ID;
    STR SSN;
    COLLECTION ORDERS_FOR_CUSTOMER;
    GET_CUST_ID ();
    SET_CUST_ID (INT CUST_ID);
    GET_SSN ();
    SET_SSN (STR SSN);
    GETORDERSFORCUSTOMER ();
    ADDORDERSFORCUSTOMER(ORDER O);
    REMOVEORDERSFORCUSTOMER(ORDER O);
}

```

424

```


CLASS ORDER {
    INT ORDER_ID;
    STR DATE;
    CUSTOMER CUSTOMER_FOR_ORDER;
    GET_ORDER_ID ();
    SET_ORDER_ID (INT ORDER_ID);
    GET_DATE ();
    SET_DATE (STR DATE);
    GETCUSTOMERFORORDER ();
    SETCUSTOMERFORORDER(CUSTOMER C);
}

```


116

FIG. 4B



5/9

**FIG. 5**

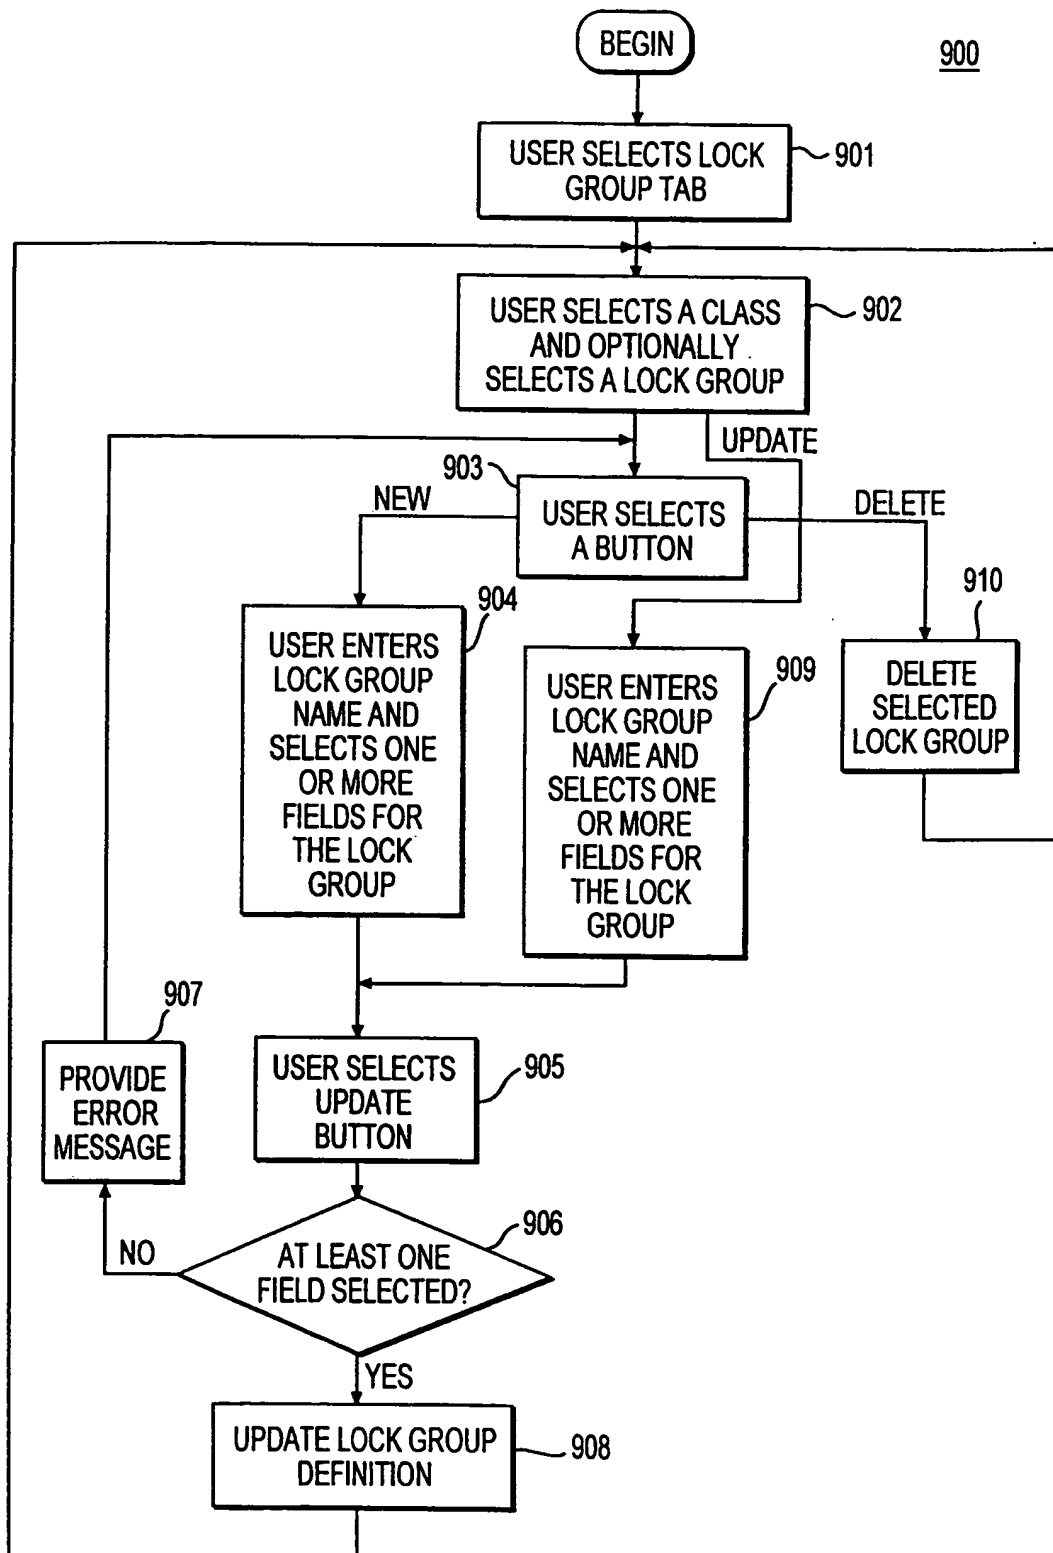
6/9

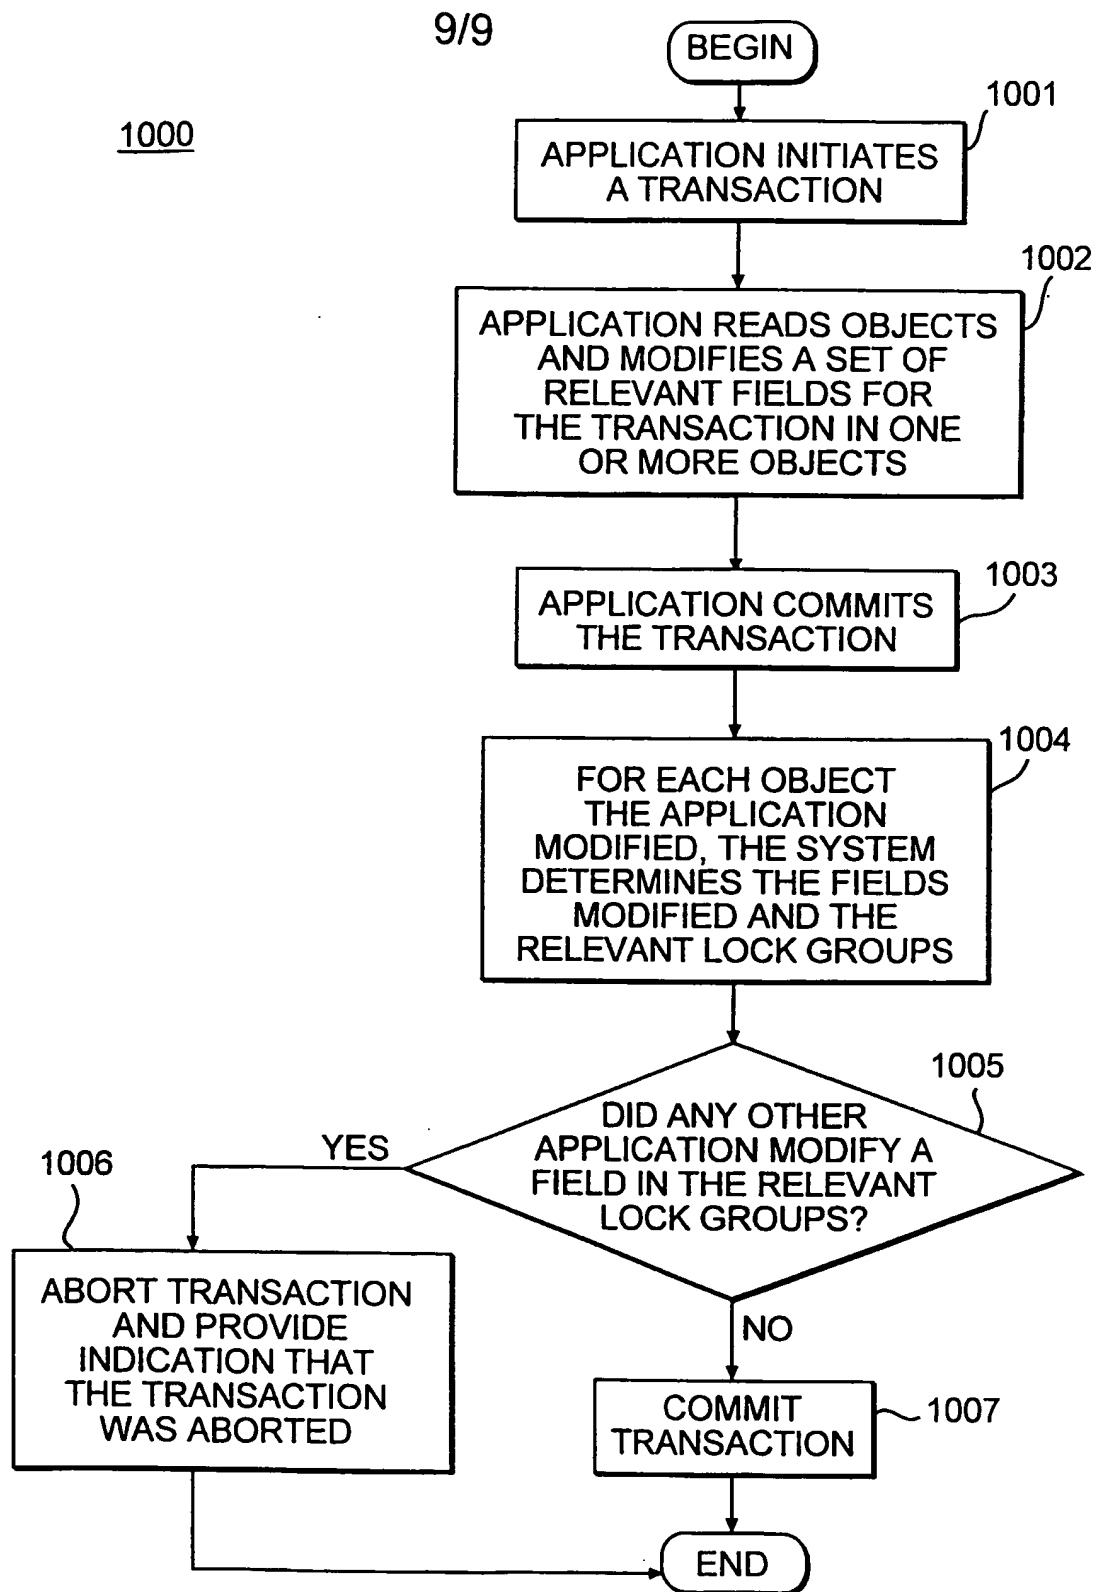


7/9

**FIG. 7****FIG. 8**

8/9

**FIG. 9**

**FIG. 10**

INTERNATIONAL SEARCH REPORT

Int lional Application No

PCT/US 98/27240

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 G06F17/30

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	"THE ENTERPRISE OBJECTS FRAMEWORK" ENTERPRISE OBJECTS FRAMEWORK: BUILDING REUSABLE BUSINESS OBJECTS, July 1994, pages 1-13, XP002047858 see page 6, left-hand column, line 1 - right-hand column, line 38; figures 3,4	1-16

☐ Further documents are listed in the continuation of box C.

☐ Patent family members are listed in annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- "&" document member of the same patent family

Date of the actual completion of the international search

21 April 1999

Date of mailing of the international search report

28/04/1999

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3018

Authorized officer

Fournier, C